



Prof. R. Wattenhofer


Generalizing Audio Source Separation



Source separation is an established research field in signal processing. Most people have probably heard of a specific kind of source separation problem, called the Cocktail Party Problem, where one hears a mixture of many voices and wants to focus in on one particular conversation.

In this thesis we want to take a look at the more general audio source separation problem and see how well we can separate with a deep learning approach any two sounds.

Requirements: Knowledge in Deep Learning, or solid background in Machine Learning. Implementation experience is an advantage. You should be able to read and understand the first 12 chapters of the "Deep Learning Book" by Goodfellow et al. (available for free online from MIT press). If you are interested in the topic but new to deep learning we expect you to complete an introductory deep learning course before applying for the thesis, such as Andrew Ng's coursera course (use the free trial!)¹ or this Udacity course².



Contacts

- Oliver Richter: richtero@ethz.ch, ETZ G63
- Gino Brunner: brunnegi@ethz.ch, ETZ G63

¹<https://www.coursera.org/specializations/deep-learning>

²<https://classroom.udacity.com/courses/ud730>